Robert M. Farber

Supercomputing for the Masses:

Introducing CUDA and GPU Computing

Friday, December 2

1:40—2:30 pm in MEC 114

Mr. Farber will remain after the presentation to continue questions and further discussion

GPU computing has made teraflop supercomputing available to anyone with a computer. Algorithm, application and library developers need to be aware of and consider the potential in GPU computing and how it now extends into conventional multi-core x86 computing. NVIDIA introduced CUDA for GPU computing in February 2007. The rate of adoption has been remarkable as have been the improvements in application performance (10-times to 1000-times) for a variety of problem domains. NVIDA estimates that over a 1/3 billion CUDA-enabled GPUs have been sold world-wide. CUDA is now taught at 454 institutions worldwide. This talk will discuss how simple it is to express problems in CUDA and particularly with the Thrust API. Results for a generic machine-learning data mining problem on a single GPU show an 85-times speedup over a modern quad-core Xeon processor (341-times single core performance) for a PCA/ NLPCA problems using Nelder-Mead. The parallel mapping developed by Farber at Los Alamos is generally applicable to a range of optimization problems (SVM, MDS, EM, ICS, ...) and optimization methods (Powell, Levenberg-Marquardt, Conjugate Gradient, ...). Scaling results will demonstrate that this same mapping, and CUDA implementation exhibits near linear scaling to 500 GPUs. A CPU version scales to over 60,000 processing cores and delivers over 1/3 of a petaflop. Speedups using CUDA in a number of other problems domains plus links to downloadable source code will be provided. Finally, recent developments make CUDA a potential development language like Java, FORTRAN, and C++ for all application development including those applications intended for only x86 architecture deployments.

Rob Farber *has served as a scientist worldwide at the Irish Center for High-End Computing, U.S. national labs in Los Alamos, Berkeley, and the Pacific Northwest, and external faculty at the Santa Fe Institute. A co-founder of two successful start-ups and a consultant to Fortune 100 companies, his articles have appeared in Dr. Dobb's Journal and Scientific Computing, among others. He recently completed his book "CUDA Application Design and Development", which is now available for purchase.*